

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: de Jesus Valdes et al.

Assignee: SRI International, Inc.

Title: Probabilistic Alert Correlation

Serial No.: 09/944,788

Filing Date: August 31, 2001

Examiner: Cristina O. Sherr

Art Unit: 3685

Mail Stop APPEAL BRIEF-PATENTS
COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, VA 22313-1450

SIR:

APPEAL BRIEF

This Brief is submitted in conjunction with the Notice of Appeal filed June 18, 2010. The Appellants request review of the rejections in the Final Office Action mailed March 18, 2010 ("Final Office Action"). Allowance of the pending claims is respectfully requested for the reasons provided below.

Table of Contents

1.	Identification Page.....	1
2.	Table of Contents	2
3.	Real Party in Interest	3
4.	Related Appeals and Interferences	3
5.	Status of Claims	3
6.	Status of Amendments	3
7.	Summary of Claimed Subject Matter	3
8.	Grounds of Rejection to be Reviewed on Appeal	7
9.	Arguments	7
10.	Conclusion	25
11.	Claims Appendix	26
12.	Evidence Appendix	44
13.	Related Proceedings Appendix	45

Real Party in Interest

The real party in interest is SRI International, Inc.

Related Appeals and Interferences

This Application is a continuation-in-part of United States Patent Application Serial No. 09/653,066, now abandoned. Application Serial No. 09/653,066 was subject to an appeal that was decided on July 13, 2009.

Appellants assert that no other appeals or interferences are known to Appellants, Appellants' legal representative, or assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

Status of Claims

Claims 1-30 are pending in the application. Claims 1-5 were originally presented in the application. Claims 3-6, 9-12, and 15-30 are withdrawn. Claims 1-2, 7-8, and 13-14 have been amended. The rejection of claims 1-2, 7-8, and 13-14 is appealed.

Status of Amendments

All claim amendments made before March 18, 2010 have been entered. No amendments were submitted subsequent to March 18, 2010.

Summary of Claimed Subject Matter

Embodiments of the present invention are generally directed to a probabilistic alert correlation method and system. In one embodiment an intrusion detection system includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected. A method for organizing these alerts into alert classes includes receiving a new alert, and identifying a set of similar features shared by the new alert and one or more existing alert classes. A threshold similarity requirement and a similarity expectation for one or more of the similar features are then updated. The new alert is next compared with the

existing alert classes and either associated with an existing alert class that it most closely matches or placed in a newly defined alert class.

For the convenience of the Board, Appellants' independent claims 1, 7, and 13 are presented below with citations to various figures and appropriate citations to at least one portion of the specification for elements of the appealed claim.

Claims 1, 7, and 13 recite (with references to illustrative portions of the specification added):

1. (Previously Presented) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected (Pg. 6, lines 12-15; Fig. 3: 301, 303), a method for organizing the alerts into alert classes, both the alerts and the alert classes having a plurality of features (Pg. 6, lines 20-23; Fig. 3: 305), the method comprising:

- (a) receiving a new alert (Pg. 6, lines 15-18; Fig. 3: 305);
- (b) identifying a set of similar features shared by the new alert and one or more existing alert classes (Pg. 10: ll. 25-28; Fig. 4: 401);
- (c) updating a threshold similarity requirement for one or more of the similar features (Pg. 10, lines 30-31; Fig. 4: 404);
- (d) updating a similarity expectation for one or more of the similar features (Pg. 10, lines 28-29; Fig. 4: 403);
- (e) comparing the new alert with the one or more existing alert classes (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 405); and either:
 - (f1) associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 407); or

(f2) defining a new alert class that is associated with the new alert (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 409),

wherein at least one of: the receiving, the identifying, the updating a threshold similarity, the updating a similarity expectation, the comparing, the associating, or the defining is performed by a processor.

7. (Previously Presented) A computer readable storage medium containing an executable program for organizing alerts that are generated by a plurality of sensors into alert classes, both the alerts and the alert classes having a plurality of features (Pg. 6, lines 20-23; Fig. 3: 305), where the program causes a processor to perform steps of:

(a) receiving a new alert (Pg. 6, lines 15-18; Fig. 3: 305);

(b) identifying a set of similar features shared by the new alert and one or more existing alert classes (Pg. 10: ll. 25-28; Fig. 4: 401);

(c) updating a threshold similarity requirement for one or more of the similar features (Pg. 10, lines 30-31; Fig. 4: 404);

(d) updating a similarity expectation for one or more of the similar features (Pg. 10, lines 28-29; Fig. 4: 403);

(e) comparing the new alert with the one or more existing alert classes (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 405); and either:

(f1) associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 407); or

(f2) defining a new alert class that is associated with the new alert (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 409).

13. (Previously Presented) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected (Pg. 6, lines 12-15; Fig. 3: 301, 303), a system for organizing the alerts into alert classes, both the alerts and the alert classes having a plurality of features (Pg. 6, lines 20-23; Fig. 3: 305), where the system comprises:

(a) means for receiving a new alert (Pg. 6, lines 15-18; Fig. 3: 305);

(b) means for identifying a set of similar features shared by the new alert and one or more existing alert classes (Pg. 10: ll. 25-28; Fig. 4: 401);

(c) means for updating a threshold similarity requirement for one or more of the similar features (Pg. 10, lines 30-31; Fig. 4: 404);

(d) means for updating a similarity expectation for one or more of the similar features (Pg. 10, lines 28-29; Fig. 4: 403);

(e) means for comparing the new alert with the one or more existing alert classes (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 405);
and

(f1) means for associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 407), or defining a new alert class that is associated with the new alert (Pg. 10, line 31 – Pg. 11, line 3; Fig. 4: 409).

Grounds of Rejection to be Reviewed on Appeal

I. Claims 1-2 are rejected under 35 U.S.C. §101 as being allegedly directed to non-statutory subject matter.

II. Claims 1-2, 7-8, and 13-14 are rejected under 35 U.S.C. §102(a) as being allegedly anticipated by Nine et al., U.S. Patent No. 6,560,611 ("Nine").

III. Claims 7-8 are rejected under 35 U.S.C. §102(a) as being allegedly anticipated by Baggen, U.S. Patent No. 4,667,317 ("Baggen").

ARGUMENTS

I. Rejection of claims 1-2 under 35 U.S.C. § 101

Claims 1-2 are rejected under 35 U.S.C. §101 as being allegedly directed to non-statutory subject matter.

A. Claim 1

Claim 1 is rejected under 35 U.S.C. § 101 as being allegedly directed to non-statutory subject matter. The Appellants respectfully disagree.

Claim 1 clearly recites that "at least one of: the receiving, the identifying, the updating a threshold similarity, the updating a similarity expectation, the comparing, the associating, or the defining is performed by a processor" (emphasis added). As such, independent claim 1 clearly recites a method that is tied to a particular machine or apparatus (*i.e.*, a processor) that performs at least one of the recited steps. Nevertheless, the Examiner submits that "the device or machine represents merely extra-solution activity, as part of a preamble" (Final Office Action, Page 4). The Appellants respectfully disagree.

First, the Appellants submit that independent claim 1 clearly satisfies the machine prong of the so-called "machine-or-transformation" test. In particular, as discussed above, at least one of the claimed steps is performed by a processor. The Appellants submit that this recitation is enough to tie the claimed invention to

a particular apparatus for the purposes of the machine-or-transformation test. For example, the Board of Patent Appeals and Interferences observed in *Ex parte Caccavale*, 2009-006026 (July 23, 2010) that the applicants' claim 8 failed to satisfy the machine-or-transformation test because it "simply fails to recite that the [claimed] computations are performed by the [claimed] 'distributed processing units' or any other machine," See *Ex parte Caccavale*, 2009-006026 at 10. Thus, the Board appears to suggest in this statement that the machine-or-transformation test can be satisfied if a claimed step (e.g., performing a computation) is explicitly recited as being performed by a machine (e.g., a distributed processing system). In the case of the Appellants' independent claim 1, at least one of the claimed steps is clearly and explicitly recited as being "performed by a processor," as discussed above (emphasis added). As such, the Appellants submit that independent claim 1 clearly satisfies the machine-or-transformation test.

The Board's post-*Bilski* decisions in *Ex parte Proudler*, 2009-006599 (July 7, 2010) and *Ex parte Bigler*, 2009-006556 (July 15, 2010) provide additional support for finding that the claimed invention satisfies the machine-or-transformation test. For example, the Board appears to suggest in *Bigler* that in order for the recitation of a processor to render a claim patent-eligible, the processor must "necessarily [be] hardware or a computer itself" and/or must be "directly claimed either in the preamble ... or the body of the claim," See *Ex parte Bigler*, 2009-006556 at 6. In the instant case, the claimed processor is necessarily a computer or part of a computing device. Specifically, the claimed processor is recited as being incorporated "[i]n an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected ...," See, e.g., Appellants' independent claim 1 (emphasis added). Sensors in this case are devices that detect intrusions and generate alerts. The processor is a computer or part of a computing device that further examines these alerts by executing the steps of the claimed method. Thus, the claimed sensors and processor are physical devices do more than merely pass signals or encompass data objects. Moreover, the Appellants note that the

sensors are directly claimed in the preamble, while the processor is directly claimed in the body of independent claim 1.

Second, the Appellants note that the machine-or-transformation test does not require that every step of a claimed method be tied to a particular machine or apparatus. As clearly recited in claim 1, at least one of the recited steps is performed by a processor. Thus, the claimed process includes at least one step that cannot be "performed by a person alerting another person via a shout, for example, or via mental steps in comparing one alert with another," as suggested by the Examiner (Final Office Action, Pages 4-5). Thus, as discussed above, the independent claim 1 clearly satisfies the machine-or-transformation test.

Third, with respect to the Examiner's allegation that the recitation of the processor in independent claim 1 represents merely extra-solution activity, the Appellants note that the "Interim Examination Instructions For Evaluating Subject Matter Eligibility Under 35 U.S.C. § 101," effective August 24, 2009, define "extra-solution activity" as "activity that is not central to the purpose of the method invented by the applicant" (Interim Examination Instructions, Page 6). The Appellants note that the amendments that were previously made to independent claim 1 in order to comply with 35 U.S.C. §101 did not include additional steps or activities. Instead, these amendments clarified the manner in which the existing steps or activities are performed (*i.e.*, in at least one case, using a processor). Moreover, the Appellants submit that all of the existing steps and activities are central to the purpose of the claimed method. Specifically, all of the steps of the claimed method are central to the purpose of organizing intrusion detection system alerts indicative of attacks or anomalous incidents into alert classes. Thus, the Appellants respectfully submit that independent claim 1 does not include "extra-solution activity."

Fourth, the Appellants note that the "device or machine" (*i.e.*, the claimed processor) is not recited as part of the preamble, as alleged by the Examiner. The claimed processor is clearly recited in the body of independent claim 1, and, as such, constitutes a positive limitation on the scope of the claim.

Fifth, as noted by the United States Supreme Court in *Bilski v. Kappos*, 561 U.S. ____ (2010), although the so-called “machine-or-transformation test” is “a useful and important clue, an investigative tool, for determining whether some claimed inventions are processes under §101,” it is “not the sole test for deciding whether an invention is a patent-eligible ‘process,’” See *Bilski v. Kappos*, 561 U.S. ____ (2010) at 8. In other words, factors beyond the machine-or-transformation test may weight toward patent eligibility. In the case of the Appellants’ independent claim 1, the claim describes a particular solution (*i.e.*, correlation of intrusion detection system alerts) to a problem to be solved (*i.e.*, efficient processing of and response to a large number of alerts). As such, the Appellants respectfully submit that independent claim 1 clearly satisfies at least this additional factor with respect to the patent eligibility of a process.

Thus, when considered as a whole, Appellants’ independent claim 1 is clearly directed to subject matter that is statutory within the meaning of 35 U.S.C. § 101.

B. Claim 2

Claim 2 is rejected under 35 U.S.C. § 101 as being allegedly directed to non-statutory subject matter. The Appellants respectfully disagree.

This ground of rejection is predicated on the validity of the rejection under 35 U.S.C. §101 as applied to independent claim 1 above. As articulated above, the Appellants submit that independent claim 1 is clearly directed to subject matter that is statutory within the meaning of 35 U.S.C. § 101. Claim 2 depends from independent claim 1 and recites at least all of the same features recited in independent claim 1. As such, the Appellants submit that the subject matter recited in claim 2 is statutory for at least the same reasons that the subject matter recited in independent claim 1 is statutory.

Thus, claim 2 is clearly directed to subject matter that is statutory within the meaning of 35 U.S.C. § 101.

II. Rejection of claims 1-2, 7-8, and 13-14 under 35 U.S.C. §102(a)

Claims 1-2, 7-8, and 13-14 are rejected under 35 U.S.C. § 102(a) as being anticipated by Nine.

The invention claimed by the Appellants allows one to correlate large numbers of intrusion detection alerts into classes based on their feature similarities before forwarding the classes to a system operator. This, in turn, allows a system operator to identify and to respond more quickly to intrusions or anomalies in a network. For instance, during a large-scale hacker attack, each sensor in an intrusion detection system may generate hundreds of alerts, and although each alert may be accurate, the sheer number can easily overwhelm a system operator. Moreover, false alarms can be triggered by normal traffic directed towards malfunctioning network resources (e.g., servers), thereby distracting the system operator from genuine intrusions or anomalies. By correlating the alerts in the manner claimed by the Appellants, genuine intrusions and anomalies can be identified, and therefore rectified, more efficiently.

A. Claim 1

Claim 1 is rejected under 35 U.S.C. §102(a) as being anticipated by Nine. The Appellants respectfully disagree.

The Appellants respectfully submit that the Examiner has failed to establish a factual basis to support the legal conclusion of anticipation. "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Examiner has failed to establish a case of anticipation because Nine fails to disclose or suggest all of the elements recited in independent claim 1.

In general, Nine discloses a system that monitors services and conditions on various networks, provides an advance warning of potential failures, and diagnoses problems. (Nine, Abstract). However, contrary to the Examiner's

suggestion, Nine does not appear to teach at least the following features recited in independent claim 1:

"comparing the new alert with the one or more existing alert classes;"

"updating a threshold similarity requirement for one or more of the similar features;"

"updating a similarity expectation for one or more of the similar features;"

"associating the new alert with the existing alert class that the new alert most closely matches;"

"defining a new alert class that is associated with the new alert;"
and

"identifying a set of similar features shared by the new alert and one or more existing alert classes."

Each of these features is discussed in detail below:

1. "Comparing the new alert with the one or more existing alert classes"

The Appellants submit that Nine is completely devoid of any teaching, showing, or suggestion relating to the comparison of an alert (indicating an attack or anomalous incident) – or more specifically, the comparison of features of the alert - to the features of existing alert classes, in order to classify the alert, as claimed by the Appellants in independent claim 1.

By contrast, Nine discloses a network monitoring system that simply reports a detected problem to the proper individual (e.g., technician), based on the nature of the problem. That is, Nine does not classify the detected problem (e.g., in accordance with its features) by comparing it to known problems, but simply evaluates the detected problem as a discrete incident and reports it to a human technician for further action.

Specifically, Nine teaches a remote monitoring system (RMS) that reports to a network operation site (NOS) when the RMS detects an anomaly with respect to a service it monitors. The report provided by the RMS is a ticket or data record containing information about the service (e.g., location, severity of problem, time of occurrence). In addition, the system "determines the nature of the problem, and notifies the proper personnel [e.g., a technician]" (See, Nine at column 3, lines 25-27). The Examiner alleges that this amounts to "a comparison of an alert in order to classify the alert" (Final Office Action, Page 3). The Appellants respectfully disagree.

The Appellants note that Nine does not explicitly disclose that the process of determining where to place a pending ticket includes a comparison or classification step. In fact, Nine is rather vague in explaining how the proper location for the ticket is determined. At best, Nine discloses an example wherein an accounting engine is "queried" for the location with the IP address and port number of a nonresponsive service. However, Nine does not disclose that "querying" involves comparing or classifying, much less what might be compared in such a case. It requires a significant leap of logic to suggest that deciding where to place a ticket is necessarily the same as classifying the pending ticket by comparing it to other pending tickets. There are many possible ways in which the proper location for a pending ticket could be determined. Nine appears to disclose determining where to place a pending ticket based solely on information contained in the pending ticket (e.g., IP address or port number; See, Nine, column 8, lines 41-43); nothing in Nine even alludes to the possibility that comparing to information contained in other tickets may be useful in determining where to place the pending ticket.

Thus, there is simply no support in Nine for the step of comparing a pending ticket to other tickets for classification purposes.

2. "Updating a threshold similarity requirement for one or more of the similar features" and "Updating a similarity expectation for one or more of the similar features"

The Appellants submit that Nine is completely devoid of any teaching, showing, or suggestion relating to "updating a threshold similarity requirement for one or more features" of an alert relative to features of alert classes or to "updating a similarity expectation for one or more features" of an alert relative to features of alert classes, as claimed by the Appellants in independent claim 1. The first portion of Nine that the Examiner cites to teach these features in fact merely teaches that the monitoring software is replicated for each service on a device by an informer engine executing forker software and sender software (See, e.g. Nine at column 5, line 45 – column 6, line 9). There is no discussion of examining the features of an alert, or of the need to update a threshold similarity requirement or a similarity expectation for the features of the alert relative to the one or more alert classes. Nine simply copies software.

The second portion of Nine that the Examiner cites to teach these features in fact merely teaches that information from the ticket file may be extracted and used to generate a report that helps to "detect patterns in problems experienced by a device" (See, e.g., Nine, column 9, lines 30-33). However, Nine does not disclose specifically how such patterns are detected. In particular, Nine does not disclose that patterns are detected by updating a threshold similarity requirement or a similarity expectation for alert features. In fact, Nine does not disclose the use of any kind of threshold or the updating of any kind of feature-related metric. As such, the Appellants respectfully submit that the Examiner is reading far more into Nine than is supported by Nine's disclosure.

Thus, there is simply no support in Nine for the steps of updating or otherwise implementing: (1) a threshold similarity requirement for one or more features of an alert relative to features of alert classes; or (2) a similarity expectation for one or more features of an alert relative to features of alert classes.

3. "Associating the new alert with the existing alert class that the new alert most closely matches"

The Appellants submit that Nine is completely devoid of any teaching, showing, or suggestion relating to associating a new alert with the existing alert class that the new alert most closely matches, as claimed by the Appellants in independent claim 1. The portion of Nine that the Examiner cites to teach this feature in fact merely teaches three techniques for detecting a problem with a monitored service. The first technique checks to make sure that the service is responsive (e.g., by "ping, nmap, finger, or telnet", Nine at column 7, lines 25-33). The second technique monitors environmental sensors to detect problems with the environment (e.g., "if the temperature is too high", Nine at column 7, lines 34-39). The third technique examines a log of the monitored service and parses for potential problems (e.g., indication that a particular route associated with a router is not functioning, Nine at column 7, lines 40-46). None of these techniques involve the comparison of an alert to existing alert classes, or the association of the alert with one of the existing alert classes based on the comparison.

Thus, there is simply no support in Nine for the step of associating a new alert with the existing alert class that the new alert most closely matches.

4. "Defining a new alert class that is associated with the new alert"

The Appellants submit that Nine is completely devoid of any teaching, showing, or suggestion relating to associating a new alert with a newly defined alert class when the new alert fails to match an existing alert class, as claimed by the Appellants in independent claim 1. The portion of Nine that the Examiner cites to teach this feature in fact merely teaches that log files for a monitored service may be used to diagnose problems with the service. Again, there is no mention of the need to compare an alert with existing alert classes in order to classify the alert, as claimed by the Appellants. Additionally, Nine discloses nothing about defining a new class when an alert fails to match an existing class. At best, Nine discloses generating a new report based on information extracted from logged tickets (such as total number of tickets). However, the generation of

the report has nothing to do with the attempted classification of a particular ticket or alert.

Thus, there is simply no support in Nine for the step of associating a new alert with a newly defined alert class when the new alert fails to match an existing alert class.

5. "Identifying a set of similar features shared by the new alert and one or more existing alert classes"

The Appellants submit that Nine is completely devoid of any teaching, showing, or suggestion relating to classifying an alert in accordance with its features, as claimed by the Appellants in independent claim 1. The portion of Nine that the Examiner cites to teach this feature in actuality merely teaches that software monitors a service and reports to the NOS when the service is unresponsive or when an anomaly is detected. The report contains "information about the service, such as location, severity of the problem, and time of occurrence" (See, e.g., Nine at column 3, lines 12-20). There is no mention in this passage of the need to identify features of the problem or to compare the problem to other known problems (e.g., existing alert classes) based on the identified features.

Thus, there is simply no support in Nine for the step of identifying a set of similar features shared by a new alert and one or more existing alert classes.

In short, as discussed above, Nine fails to teach, show, or suggest any sort of classification of alerts by comparing features of the alerts to features of existing alert classes, as recited by the Appellants in independent claim 1. Moreover, the Appellants respectfully submit that the explicit teachings of Nine actually teach away from the claimed classification step. Specifically, Nine teaches that an additional, post-processing reporting feature is required to detect groups of tickets related to a common problem (See, e.g., Nine, column 9, lines 30-39: "if a series of tickets indicate that a security log file on an NT server has a flood of ICMP packets, a report may be created to locate all of the tickets that

indicate this problem,” emphasis added). If the tickets had been classified (*i.e.*, compared against other tickets and grouped together into classes) as they were generated (*i.e.*, before being transmitted to the appropriate location), then such a report would not be necessary. That is, all of the tickets that indicate the problem would already be grouped together. Thus, the post-transmission reporting feature required by Nine clearly indicates that Nine teaches that the tickets are not classified or compared to each other, as claimed by the Appellants.

As discussed above, correlating the alerts in the manner claimed by the Appellants allows a system operator to identify genuine intrusions and anomalies and therefore rectify these intrusions and anomalies more efficiently. As discussed above, Nine fails to disclose or suggest several of the claimed steps in this process. Thus, Nine fails to anticipate the Appellants’ independent claim 1.

Therefore, Appellants’ independent claim 1 is patentable under 35 U.S.C. §102(a) over Nine.

B. Claim 2

Claim 2 is rejected under 35 U.S.C. §102(a) as being anticipated by Nine. The Appellants respectfully disagree.

This ground of rejection is predicated on the validity of the rejection under 35 U.S.C. §102 given Nine as applied to independent claim 1 above. As articulated above with respect to independent claim 1, there are features recited in independent claim 1 that are not disclosed or suggested by Nine, including:

“comparing the new alert with the one or more existing alert classes;”

“updating a threshold similarity requirement for one or more of the similar features;”

“updating a similarity expectation for one or more of the similar features;”

“associating the new alert with the existing alert class that the new alert most closely matches;”

"defining a new alert class that is associated with the new alert;"
and

"identifying a set of similar features shared by the new alert and
one or more existing alert classes."

Thus, dependent claim 2, which depends from independent claim 1 and recites at least all of the same features recited in independent claim 1, has been erroneously rejected under 35 U.S.C. §102(a). The Examiner has failed to establish a showing of anticipation.

In addition, Nine fails to disclose "passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes," as recited by claim 2. By contrast, the first cited portion of Nine (*i.e.*, column 5, line 60 – column 6, line 10) merely discloses how the forker software works. Specifically, the forker software requests a list of all services on a device to be monitored, and informer software responds with the port number and address of the services. Nine does not disclose that this process relies on or requires a transition model or a belief state, and the Appellants fail to see what purpose a transition model or belief state would serve in this instance.

The second cited portion of Nine (*i.e.*, column 9, lines 22-40) discloses that reporter software generates reports by auditing and extracting information from the central accounting system (CAS) database. These reports may then be used to detect patterns in problems experienced by a device, which may in turn lead to detection of larger-scale problems. Again, however, Nine does not disclose that this process relies on or requires a transition model or a belief state, and the Appellants fail to see what purpose a transition model or belief state would serve in this instance. In fact, Nine does not disclose specifically how reports may be created to detect problems, just that they may be created.

Therefore, Appellants' claim 2 is patentable under 35 U.S.C. §102(a) over Nine.

C. Claim 7

Claim 7 is rejected under 35 U.S.C. §102(a) as being anticipated by Nine. The Appellants respectfully disagree.

As discussed in detail above, Nine does not appear to disclose or suggest at least the following features recited in independent claim 7:

"comparing the new alert with the one or more existing alert classes;"

"updating a threshold similarity requirement for one or more of the similar features;"

"updating a similarity expectation for one or more of the similar features;"

"associating the new alert with the existing alert class that the new alert most closely matches;"

"defining a new alert class that is associated with the new alert;"
and

"identifying a set of similar features shared by the new alert and one or more existing alert classes."

Therefore, Appellants' independent claim 7 is patentable under 35 U.S.C. §102(a) over Nine.

D. Claim 8

Claim 8 is rejected under 35 U.S.C. §102(a) as being anticipated by Nine. Appellants respectfully disagree.

This ground of rejection is predicated on the validity of the rejection under 35 U.S.C. §102 given Nine as applied to independent claim 7 above. As articulated above with respect to independent claim 7, there are features recited in independent claim 7 that are not disclosed or suggested by Nine, including:

"comparing the new alert with the one or more existing alert classes;"

"updating a threshold similarity requirement for one or more of the similar features;"

"updating a similarity expectation for one or more of the similar features;"

"associating the new alert with the existing alert class that the new alert most closely matches;"

"defining a new alert class that is associated with the new alert;"
and

"identifying a set of similar features shared by the new alert and one or more existing alert classes."

Thus, dependent claim 8, which depends from independent claim 7 and recites at least all of the same features recited in independent claim 7, has been erroneously rejected under 35 U.S.C. §102(a). The Examiner has failed to establish a showing of anticipation.

In addition, as discussed in detail above, Nine fails to disclose "passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes," as recited by claim 8.

Therefore, Appellants' claim 8 is patentable under 35 U.S.C. §102(a) over Nine.

E. Claim 13

Claim 13 is rejected under 35 U.S.C. §102(a) as being anticipated by Nine. Appellants respectfully disagree.

As discussed in detail above, Nine does not appear to disclose or suggest at least the following features recited in independent claim 13:

"means for comparing the new alert with the one or more existing alert classes;"

"means for updating a threshold similarity requirement for one or more of the similar features;"

"means for updating a similarity expectation for one or more of the similar features;"

"means for associating the new alert with the existing alert class that the new alert most closely matches;"

"means for defining a new alert class that is associated with the new alert;" and

"means for identifying a set of similar features shared by the new alert and one or more existing alert classes."

Therefore, Appellants' independent claim 13 is patentable under 35 U.S.C. §102(a) over Nine.

F. Claim 14

Claim 14 is rejected under 35 U.S.C. §102(a) as being anticipated by Nine. Appellants respectfully disagree.

This ground of rejection is predicated on the validity of the rejection under 35 U.S.C. §102 given Nine as applied to independent claim 13 above. As articulated above with respect to independent claim 13, there are features recited in independent claim 13 that are not disclosed or suggested by Nine, including:

"comparing the new alert with the one or more existing alert classes;"

"updating a threshold similarity requirement for one or more of the similar features;"

"updating a similarity expectation for one or more of the similar features;"

"associating the new alert with the existing alert class that the new alert most closely matches;"

"defining a new alert class that is associated with the new alert;"
and

"identifying a set of similar features shared by the new alert and one or more existing alert classes."

Thus, dependent claim 14, which depends from independent claim 13 and recites at least all of the same features recited in independent claim 13, has been erroneously rejected under 35 U.S.C. §102(a). The Examiner has failed to establish a showing of anticipation.

In addition, as discussed in detail above, Nine fails to disclose "passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes," as recited by claim 14.

Therefore, Appellants' claim 14 is patentable under 35 U.S.C. §102(a) over Nine.

III. Rejection of claims 7-8 under 35 U.S.C. §102(a)

Claims 7-8 are rejected under 35 U.S.C. § 102(a) as being anticipated by Baggen.

A. Claim 7

Claim 7 is rejected under 35 U.S.C. §102(a) as being anticipated by Baggen. The Appellants respectfully disagree.

The Appellants respectfully submit that the Examiner has failed to establish a factual basis to support the legal conclusion of anticipation. "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). The Examiner has failed to establish a case of anticipation because Baggen fails to teach or suggest all of the elements of independent claim 7.

Specifically, Baggen does not appear to disclose or suggest at least the following features recited in independent claim 7:

"receiving a new alert;"

"identifying a set of similar features shared by the new alert and one or more existing alert classes;"

"comparing the new alert with the one or more existing alert classes;"

"updating a threshold similarity requirement for one or more of the similar features;"

"updating a similarity expectation for one or more of the similar features;"

"associating the new alert with the existing alert class that the new alert most closely matches;" and

"defining a new alert class that is associated with the new alert."

By contrast, Baggen is directed to a method for storing and reproducing data using a standard compact disk digital audio player (See, Baggen, Abstract). In other words, Baggen has absolutely nothing to do with intrusion detection systems, or the correlation of alerts produced by the sensors of an intrusion detection system. Thus, Baggen completely fails to disclose or suggest any of the features recited in the body of independent claim 7.

Moreover, independent claim 7 was previously amended, almost verbatim in accordance with the Examiner's suggestion, to recite that the claimed program "causes a processor to perform steps of: receiving ... identifying ... updating a threshold similarity ... updating a similarity expectation ... comparing ... associating ... and defining ..." (emphasis added). The Examiner indicated in the Final Office Action that such an amendment "would make the claims distinguishable from a generic computer readable medium with data," as allegedly disclosed by Baggen (Final Office Action, Page 7). The Advisory Action indicates that this amendment was entered. Nevertheless, the rejection under 35 U.S.C. §102 over Baggen appears to have been maintained.

Therefore, Appellants' independent claim 7 is patentable under 35 U.S.C.

§102(a) over Baggen.

B. Claim 8

Claim 8 is rejected under 35 U.S.C. §102(a) as being anticipated by Baggen. Appellants respectfully disagree.

This ground of rejection is predicated on the validity of the rejection under 35 U.S.C. §102 given Baggen as applied to independent claim 7 above. As articulated above with respect to independent claim 7, there are features recited in independent claim 7 that are not disclosed or suggested by Baggen, including:

"comparing the new alert with the one or more existing alert classes;"

"updating a threshold similarity requirement for one or more of the similar features;"

"updating a similarity expectation for one or more of the similar features;"

"associating the new alert with the existing alert class that the new alert most closely matches;"

"defining a new alert class that is associated with the new alert;"
and

"identifying a set of similar features shared by the new alert and one or more existing alert classes."

Thus, dependent claim 8, which depends from independent claim 7 and recites at least all of the same features recited in independent claim 7, has been erroneously rejected under 35 U.S.C. §102(a). The Examiner has failed to establish a showing of anticipation.

In addition, Baggen fails to disclose "passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes," as recited by claim 8. As discussed above, Baggen is directed to the storage and reproduction of data,

and not to the field of intrusion detection. Moreover, nothing in Baggen discloses or suggests that a transition model or a belief state would be useful in the storage or reproduction of data.

Therefore, Appellants' claim 8 is patentable under 35 U.S.C. §102(a) over Baggen.

IV. Conclusion

In view of the above, the Appellants submit that the rejections are improper, the claimed invention is patentable, the rejections of claims 1-2, 7-8, and 13-14 should be reversed, and the application should be allowed.

Respectfully submitted,

Dated: 8/18/10



Kin-Wah Tong, Attorney
Reg. No. 39,400
(732) 542-2280

Wall & Tong, LLP
25 James Way
Eatontown, New Jersey 07724

CLAIMS APPENDIX

1. (Previously Presented) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a method for organizing the alerts into alert classes, both the alerts and the alert classes having a plurality of features, the method comprising:

(a) receiving a new alert;

(b) identifying a set of similar features shared by the new alert and one or more existing alert classes;

(c) updating a threshold similarity requirement for one or more of the similar features;

(d) updating a similarity expectation for one or more of the similar features;

(e) comparing the new alert with the one or more existing alert classes;

and either:

(f1) associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches; or

(f2) defining a new alert class that is associated with the new alert,

wherein at least one of: the receiving, the identifying, the updating a threshold similarity, the updating a similarity expectation, the comparing, the associating, or the defining is performed by a processor.

2. (Previously Presented) The method of claim 1 further comprising a step

(a1) of passing each of the one or more existing alert classes through a transition

model to generate a new prior belief state for each of the one or more existing alert classes.

3. (Withdrawn) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a method for organizing alerts having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;

(b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) comparing the new alert to one or more alert classes;

(d) rejecting a match if any feature for which a minimum similarity value has been set fails to meet or exceed the minimum similarity value;

(e) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

(f1) associating the new alert with the existing alert class that the new alert most closely matches; or

(f2) defining a new alert class that is associated with the new alert.

4. (Withdrawn) In an intrusion detection system that includes a plurality of sensors, each of which generates alerts when attacks or anomalous incidents are detected, a method for organizing the alerts comprising the steps of:

- (a) receiving an alert;
- (b) identifying a set of features that may be shared by the received alert and one or more existing alert classes;
- (c) setting a minimum similarity value for one or more features or feature groups; comparing the new alert to one or more of the alert classes, and either:
 - (d1) defining a new alert class that is associated with the received alert if any feature or feature group that has a minimum similarity value fails to meet or exceed its minimum similarity value; or
 - (d2) associating the received alert with the existing alert class that the received alert most closely matches.

5. (Withdrawn) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a method for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, the method comprising the steps of:

- (a) receiving a new alert;
- (b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes;
- (c) updating a minimum similarity requirement for one or more features;
- (d) comparing the new alert with one or more alert classes, and either:
 - (e1) associating the new alert with the existing alert class that the new alert most closely matches; or
 - (e2) defining a new alert class that is associated with the new alert.

6. (Withdrawn) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a method for organizing alerts having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;

(b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) comparing the new alert to one or more alert classes;

(d) rejecting a match if any feature for which a minimum similarity value has been set fails to meet or exceed the minimum similarity value, and either:

(e1) associating the new alert with the existing alert class that the new alert most closely matches; or

(e2) defining a new alert class that is associated with the new alert.

7. (Previously Presented) A computer readable storage medium containing an executable program for organizing alerts that are generated by a plurality of sensors into alert classes, both the alerts and the alert classes having a plurality of features, where the program causes a processor to perform steps of:

(a) receiving a new alert;

(b) identifying a set of similar features shared by the new alert and one or more existing alert classes;

(c) updating a threshold similarity requirement for one or more of the similar features;

(d) updating a similarity expectation for one or more of the similar features;

(e) comparing the new alert with the one or more existing alert classes;

and either:

(f1) associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches; or

(f2) defining a new alert class that is associated with the new alert.

8. (Previously Presented) The computer readable storage medium of claim 7 further comprising a step (a1) of passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes.

9. (Withdrawn) A computer readable medium containing an executable program for organizing alerts that are generated by a plurality of sensors and have a plurality of features, each feature having one or more values, where the program performs the steps of:

(a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;

(b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

- (c) comparing the new alert to one or more alert classes;
- (d) rejecting a match if any feature for which a minimum similarity value has been set fails to meet or exceed the minimum similarity value;
- (e) adjusting the comparison by an expectation that certain feature values will or will not match, and either:
 - (f1) associating the new alert with the existing alert class that the new alert most closely matches; or
 - (f2) defining a new alert class that is associated with the new alert.

10. (Withdrawn) A computer readable medium containing an executable program for organizing alerts generated by a plurality of sensors, where the program performs the steps of:

- (a) receiving an alert;
- (b) identifying a set of features that may be shared by the received alert and one or more existing alert classes;
- (c) setting a minimum similarity value for one or more features or feature groups; comparing the new alert to one or more of the alert classes, and either:
 - (d1) defining a new alert class that is associated with the received alert if any feature or feature group that has a minimum similarity value fails to meet or exceed its minimum similarity value; or
 - (d2) associating the received alert with the existing alert class that the received alert most closely matches.

11. (Withdrawn) A computer readable medium containing an executable program for organizing alerts generated by a plurality of sensors into alert classes, both the alerts and alert classes having a plurality of features, where the program performs the steps:

- (a) receiving a new alert;
- (b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes;
- (c) updating a minimum similarity requirement for one or more features;
- (d) comparing the new alert with one or more alert classes, and either:
 - (e1) associating the new alert with the existing alert class that the new alert most closely matches; or
 - (e2) defining a new alert class that is associated with the new alert.

12. (Withdrawn) A computer readable medium containing an executable program for organizing alerts generated by a plurality of sensors and having a plurality of features, each feature having one or more values, where the program performs the steps of:

- (a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;
- (b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;
- (c) comparing the new alert to one or more alert classes;
- (d) rejecting a match if any feature for which a minimum similarity value

has been set fails to meet or exceed the minimum similarity value, and either:

(e1) associating the new alert with the existing alert class that the new alert most closely matches; or

(e2) defining a new alert class that is associated with the new alert.

13. (Previously Presented) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a system for organizing the alerts into alert classes, both the alerts and the alert classes having a plurality of features, where the system comprises:

(a) means for receiving a new alert;

(b) means for identifying a set of similar features shared by the new alert and one or more existing alert classes;

(c) means for updating a threshold similarity requirement for one or more of the similar features;

(d) means for updating a similarity expectation for one or more of the similar features;

(e) means for comparing the new alert with the one or more existing alert classes; and

(f1) means for associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

14. (Previously Presented) The system of claim 13 further comprising (a1) means for passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes.

15. (Withdrawn) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a system for organizing alerts having a plurality of features, each feature having one or more values, the system comprising:

(a) means for generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;

(b) means for identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) means for comparing the new alert to one or more alert classes;

(d) means for rejecting a match if any feature for which a minimum similarity value has been set fails to meet or exceed the minimum similarity value;

(e) means for adjusting the comparison by an expectation that certain feature values will or will not match; and

(f1) means for associating the new alert with the existing alert class that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

16. (Withdrawn) In an intrusion detection system that includes a plurality of sensors, each of which generates alerts when attacks or anomalous incidents are detected, a system for organizing the alerts, the system comprising:

(a) means for receiving an alert;

(b) means for identifying a set of features that may be shared by the received alert and one or more existing alert classes;

(c) means for setting a minimum similarity value for one or more features or feature groups; comparing the new alert to one or more of the alert classes; and

(d1) means for defining a new alert class that is associated with the received alert if any feature or feature group that has a minimum similarity value fails to meet or exceed its minimum similarity value, or associating the received alert with the existing alert class that the received alert most closely matches.

17. (Withdrawn) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a system for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, the system comprising:

(a) means for receiving a new alert;

(b) means for identifying a set of potentially similar features shared by the new alert and one or more existing alert classes;

(c) means for updating a minimum similarity requirement for one or more

features;

(d) means for comparing the new alert with one or more alert classes; and

(e1) means for associating the new alert with the existing alert class that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

18. (Withdrawn) In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a system for organizing alerts having a plurality of features, each feature having one or more values, the system comprising:

(a) means for generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding feature;

(b) means for identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) means for comparing the new alert to one or more alert classes;

(d) means for rejecting a match if any feature for which a minimum similarity value has been set fails to meet or exceed the minimum similarity value; and

(e1) means for associating the new alert with the existing alert class that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

19. (Withdrawn) A method for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) identifying a set of potentially similar features shared by a new alert and one or more existing alert classes;

(b) comparing the new alert to one or more existing alert classes;

(c) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

(d1) associating the new alert with the existing alert class that the new alert most closely matches; or

(d2) defining a new alert class that is associated with the new alert.

20. (Withdrawn) A method for organizing alerts into alert classes, both the alerts and the alert classes having a plurality of features, each of the plurality of features having one or more values, the method comprising the steps of:

(a) receiving a new alert;

(b) identifying a set of similar features shared by the new alert and one or more existing alert classes;

(c) updating a similarity expectation for one or more feature values;

(d) comparing the new alert with the one or more existing alert classes;

and either:

(e1) associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches; or

(e2) defining a new alert class that is associated with the new alert.

21. (Withdrawn) The method of claim 20 further comprising the step (a1) of passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes.

22. (Withdrawn) A method for organizing alerts having a plurality of features, each feature having one or more values, the method comprising the steps of:

(a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding features;

(b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) comparing the new alert to one or more alert classes;

(d) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

(e1) associating the new alert with the existing alert class that the new alert most closely matches; or

(e2) defining a new alert class that is associated with the new alert.

23. (Withdrawn) A computer readable medium containing an executable program for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, each feature having one or more values, where the program performs the steps of:

(a) identifying a set of potentially similar features shared by a new alert and one or more existing alert classes;

(b) comparing the new alert to one or more existing alert classes;

(c) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

(d1) associating the new alert with the existing alert class that the new alert most closely matches; or

(d2) defining a new alert class that is associated with the new alert.

24. (Withdrawn) A computer readable medium containing an executable program for organizing alerts into alert classes, both the alerts and the alert classes having a plurality of features, each of the plurality of features having one or more values, where the program performs the steps of:

(a) receiving a new alert;

(b) identifying a set of similar features shared by the new alert and one or more existing alert classes;

(c) updating a similarity expectation for one or more feature values;

(d) comparing the new alert with the one or more existing alert classes;

and either:

(e1) associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches; or

(e2) defining a new alert class that is associated with the new alert.

25. (Withdrawn) The computer readable medium of claim 24 further comprising the step (a1) of passing each of the one or more existing alert classes

through a transition model to generate a new prior belief state for each of the one or more existing alert classes.

26. (Withdrawn) A computer readable medium containing an executable program for organizing alerts having a plurality of features, each feature having one or more values, where the program performs the steps of:

- (a) generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding features;

- (b) identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

- (c) comparing the new alert to one or more alert classes;

- (d) adjusting the comparison by an expectation that certain feature values will or will not match, and either:

- (e1) associating the new alert with the existing alert class that the new alert most closely matches; or

- (e2) defining a new alert class that is associated with the new alert.

27. (Withdrawn) A system for organizing alerts into alert classes, both the alerts and alert classes having a plurality of features, each feature having one or more values, the system comprising:

- (a) means for identifying a set of potentially similar features shared by a new alert and one or more existing alert classes;

(b) means for comparing the new alert to one or more existing alert classes;

(c) means for adjusting the comparison by an expectation that certain feature values will or will not match; and

(d1) means for associating the new alert with the existing alert class that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

28. (Withdrawn) A system for organizing alerts into alert classes, both the alerts and the alert classes having a plurality of features, each of the plurality of features having one or more values, the system comprising:

(a) means for receiving a new alert;

(b) means for identifying a set of similar features shared by the new alert and one or more existing alert classes;

(c) means for updating a similarity expectation for one or more feature values;

(d) means for comparing the new alert with the one or more existing alert classes; and

(e1) means for associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

29. (Withdrawn) The system of claim 28 further comprising (a1) means for passing each of the one or more existing alert classes through a transition model to generate a new prior belief state for each of the one or more existing alert classes.

30. (Withdrawn) A system for organizing alerts having a plurality of features, each feature having one or more values, the system comprising:

(a) means for generating a group of feature records for a new alert, each feature record including a list of observed values for its corresponding features;

(b) means for identifying a set of potentially similar features shared by the new alert and one or more existing alert classes that are associated with previous alerts;

(c) means for comparing the new alert to one or more alert classes;

(d) means for adjusting the comparison by an expectation that certain feature values will or will not match; and

(e1) means for associating the new alert with the existing alert class that the new alert most closely matches, or defining a new alert class that is associated with the new alert.

EVIDENCE APPENDIX

None.

10454-022001/P-4190-4 (SRI/4190-4)
09/944,788

PATENT
Conf. No. 1821



RELATED PROCEEDINGS APPENDIX

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/653,066	09/01/2000	Alfonso de Jesus Valdes	SRI/4190-2	5430
<div>52197 7590 07/13/2009 Wall & Tong, LLP SRI INTERNATIONAL 595 SHREWSBURY AVENUE SHREWSBURY, NJ 07702</div> <div>EXAMINER SIMITOSKI, MICHAEL J</div> <div>ART UNIT PAPER NUMBER 2439</div> <div>MAIL DATE DELIVERY MODE 07/13/2009 PAPER</div>				

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte ALFONSO DE JESUS VALDES,
KEITH MICHAEL SKINNER,
and PHILLIP A. PORRAS

Appeal 2008-003088
Application 09/653,066¹
Technology Center 2400

Decided:² July 13, 2009

Before ALLEN R. MacDONALD, *Vice Chief Administrative Patent Judge*,
and LEE E. BARRETT and STEPHEN C. SIU, *Administrative Patent
Judges*.

BARRETT, *Administrative Patent Judge*.

DECISION ON APPEAL

¹ Filed September 1, 2000, titled "Method for Detecting and Diagnosing Abnormalities Using Real-Time Bayes Networks.

² The two-month time period for filing an appeal or commencing a civil action, as recited in 37 C.F.R. § 1.304, begins to run from the decided date shown on this page of the decision. The time period does not run from the Mail Date (paper delivery) or Notification Date (electronic delivery).

This is a decision on appeal under 35 U.S.C. § 134(a) from the final rejection of claims 1-24. We have jurisdiction pursuant to 35 U.S.C. § 6(b).

We affirm.

STATEMENT OF THE CASE

The invention

The invention relates to detection and diagnosis of abnormalities in a computer network using real-time Bayes networks. Spec. 1. "Standard Bayes networks are probabilistic analysis tools that include fixed models (called 'hypotheses') that represent some condition or state; probabilistic reasoning methods are used to determine a belief that a current observation corresponds to one of the stored hypotheses." Spec. 3.

The claims

Illustrative claim 1 is reproduced below:

1. A method for detecting potentially harmful traffic received by a computer system or network, comprising the steps of:

(a) receiving, at a first computing device, network traffic sent between at least a second computing device and a third computing device in said network;

(b) examining the received network traffic for the presence of one or more relevant features; and

(c) providing a summary or list of relevant features contained in the received network traffic to a real-time Bayes network for analysis, the Bayes network including a model of normal traffic and a model of potentially harmful traffic as Bayes hypotheses.

The references

Toyama

US 6,502,082 B1

Dec. 31, 2002
(filed Oct. 12, 1999)

Appeal 2008-003088
Application 09/653,066

Debra Anderson et al., *Next Generation Intrusion Detection Expert Systems (NIDES) Software Users Manual*, Dec. 1, 1994, pp. 1-301 ("Anderson").

Dan Decasper et al., *Router Plugins: A Software Architecture for Next Generation Routers*, ACM, 1998, pp. 229-240 ("Decasper").

William DuMouchel, *Computer Intrusion Detection Based on Bayes Factors for Comparing Command Transition Probabilities*, Technical Report Number 91, National Institute of Statistical Sciences, Feb. 1999 ("DuMouchel").

Thomas D. Garvey and Teresa F. Lunt, *Model-Based Intrusion Detection*, 14th National Computer Security Conference, Washington, DC, Oct. 1-4, 1991, pp. 372-385 ("Garvey").

Intrusion Detection Systems Buyer's Guide, ICSA Labs, Dec. 1999, pp. 1-52 ("ICSA").

Ulf Lindqvist and Phillip A. Porras, *Detecting Computer and Network Misuse Through the Production-Based Expert System Toolset (P-BEST)*, Proc. of 1999 IEEE Symposium on Security and Privacy, Oakland, CA, May 9-12, 1999 ("Lindqvist").

Judea Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Rev. 2d printing Morgan Kaufmann Publishers, Inc. 1988), pp. 381-382 ("Pearl").

Appeal 2008-003088
Application 09/653,066

The rejections

Claims 1, 2, and 22-24 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Anderson, ICSA, and DuMouchel.

Claims 3-9, 13, 14, and 17 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Anderson, ICSA, DuMouchel, and Decasper.

Claims 10-12 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Anderson, ICSA, DuMouchel, and Decasper, further in view of Lindqvist.

Claim 15 stands rejected under 35 U.S.C. § 103(a) as unpatentable over Anderson, ICSA, DuMouchel, and Decasper, further in view of Toyama.

Claim 16 stands rejected under 35 U.S.C. § 103(a) as unpatentable over Anderson, ICSA, DuMouchel, and Decasper, further in view of Pearl.

Claims 18-20 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Anderson, ICSA, DuMouchel, and Decasper, further in view of Garvey.

PRINCIPLES OF LAW

A rejection under 35 U.S.C. § 103(a) must be based on the following factual determinations: (1) the scope and content of the prior art; (2) the level of ordinary skill in the art; (3) the differences between the claimed invention and the prior art; and (4) any objective indicia of non-obviousness.

DyStar Textilfarben GmbH & Co. Deutschland KG v. C.H. Patrick Co., 464 F.3d 1356, 1360 (Fed. Cir. 2006) (citing *Graham v. John Deere Co.*, 383 U.S. 1, 17 (1966)).

The Supreme Court states that "the [obviousness] analysis need not seek out precise teachings directed to the specific subject matter of the

Appeal 2008-003088
Application 09/653,066

challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ." *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 418 (2007). "A person of ordinary skill is also a person of ordinary creativity, not an automaton." *Id.* at 421.

DISCUSSION

Anderson, ICSA, and DuMouchel

Claims 1, 2, 22, and 23

Issue

The issue is: Would one of ordinary skill in the intrusion detection system (IDS) art would have been motivated to apply a real-time Bayes network, as taught by DuMouchel, to analysis of statistical network traffic as taught by ICSA. This is consistent with Appellants' argument that "*Anderson, ICSA and DuMouchel*, singly or in any permissible combination, fail to teach, show or suggest a network intrusion detection system that provides a summary or list of relevant features of network traffic (*i.e.*, packets or data traffic sent between network users external to a location where intrusion detection occurs) to a Bayes network for analysis, as claimed by the Appellant in independent claims 1, 2, 22 and 23" (Br. 9).

Facts

Anderson

Anderson describes the Next Generation Intrusion Detection Expert System (NIDES) (Title). "NIDES is a comprehensive intrusion-detection system that performs real-time monitoring of user activity on one or more target system computers." P. 3. "The NIDES host monitors usage on a number of computers connected via an Ethernet network." *Id.*

Anderson describes: "The agen process runs on each target host system that is actively providing audit data to NIDES for analysis. The agen program reads native audit record data, converts it into NIDES audit records, and delivers these records to the arpool process." P. 4. "The NIDES arpool process collects audit data from the various target hosts and provides the data to the analysis components for analysis and anomaly detection." P. 5.

Anderson describes: "The NIDES statistical analysis component maintains historical statistical profiles for each user and raises an alarm when observed activity departs from established patterns of use for an individual. The historical profiles are updated regularly, and older data 'aged' out with each profile update, so that NIDES adaptively learns what to expect from each user." P. 5.

Anderson describes: "The NIDES rulebased analysis component uses rules that characterize known intrusion types to raise an alarm if observed activity matches any of its encoded rules." P. 5.

Anderson also describes: "The NIDES primary mode of operation is to analyze data and report suspicious activity in real time." P. 7.

ICSA

ICSA is a guide to intrusion detection systems (IDSs).

ICSA describes five different categories of IDS: (1) network based IDS; (2) host based IDS; (3) file integrity checker; (4) network vulnerability scanner; and (5) host vulnerability scanner. P. 13. Not all of these categories represent "classical intrusion detection." *Id.*

ICSA describes:

Intrusion detection systems are either network-based or host-based; vendors are only beginning to integrate the two technologies.

Network based intrusion detection systems are most common, and examine passing traffic for signs of intrusion (See Figure 3). Host-based systems look at user and process activity on the local machine for signs of intrusion.

P. 14.

ICSA describes that IDSs generally use three kinds of available analysis engines: event or signature based, statistical analysis, and adaptive systems (p. 15). "A statistical analysis system builds statistical models of the environment, such as the average length of a telnet session, then looks for deviations from 'normal.'" P. 15.

ICSA describes that:

The network IDS usually has two logical components: the sensor and the management station. The sensor sits on a network segment, monitoring it for suspicious traffic. The management station receives alarms from the sensor(s) and displays them to an operator.

P. 15. The sensors capture passing network traffic on a segment for analysis, and "[i]f they detect something that looks unusual, they pass it back to the analysis station" (*id.*). "The analysis station can display the alarms or do additional analysis." *Id.*

DuMouchel

DuMouchel describes a statistical computer intrusion detection method based on Bayes factors.

DuMouchel describes

In computer intrusion detection one attempts to identify unauthorized accesses to computer accounts. There are two main approaches to intrusion detection: pattern recognition and anomaly detection. Pattern recognition is the attempt to recognize general patterns in command usage that stem from known attacks such as exploiting a software bug. . . . Anomaly detection, on the other hand, attempts to identify an unauthorized user by identifying unusual, for the account holder, usage of the computer. Usually, for each user a historical profile is compiled and large deviations from the profile indicates a possible intruder. Therefore it is also referred to as the profile based approach. Intrusion detection systems like IDES . . . , NIDES . . . and Emerald . . . use both approaches, presumably because neither one is uniformly superior to the other. In this paper we only consider the anomaly detection approach. This approach lends itself to a statistical treatment. Ryan et al. (1998) suggested each that each user on a computer system leaves a "print" that could be captured by training a neural network with historical data. When for new data from any user the neural network predicts that the data is more likely to stem from another user in the historical data, then an alarm for a possible intrusion is raised. . . . In this paper we propose a test for anomaly detection based on Bayesian hypothesis testing.

First page (unnumbered), "Introduction."

DuMouchel describes:

This statistical method can compare in real time the sequence of commands given by each user to a profile of that user's past behavior. We use a Bayes Factor statistic to test the null hypothesis that the observed command transition probabilities come from a profiled transition matrix.

First page (unnumbered), "Summary."

DuMouchel describes:

The null hypothesis H_0 assumes that the legitimate user has generated the data from the profiled transition probabilities. The alternative hypothesis H_1 assumes that the T commands have been drawn independently using an arbitrary probability vector.

Top of sixth page (unnumbered).

DuMouchel states that "Figures 1-3 could be constructed for virtually any intrusion detection indicator" (last page).

Contentions

The rejection, as stated, is a little confusing. The Examiner finds that Anderson discloses providing network traffic to an analysis network including a model of normal traffic and potentially harmful traffic/statistical analysis component, but does not teach examining network traffic for the presence of relevant features or the analysis component being a Bayes network (2d Final Rej. 5 entered Feb. 17, 2006; Ans. 3); i.e., the Examiner finds that Anderson teaches a network-based IDS. Later, the Examiner appears to agree with Appellants that Anderson, in fact, teaches a host-based IDS, instead of a network-based IDS (Ans. 11-12). The Examiner finds that ICSA teaches a network-based IDS that detects unusual relevant traffic and concludes that it would have been obvious to modify Anderson to also monitor relevant network traffic as taught by ICSA (2d Final Rej. 5-6; Ans. 3-4). The Examiner finds that Anderson and ICSA do not teach intrusion detection analysis with a Bayes network. The Examiner finds that DuMouchel teaches performing intrusion detection using a Bayes network that results in a low false alarm rate and high rate of expected detection and concludes that it would have been obvious to use a Bayes network in

Appeal 2008-003088
Application 09/653,066

Anderson as modified to achieve the low false alarm rate and high rate of expected detection taught by DuMouchel (2d Final Rej. 6; Ans. 4).

Appellants argue that the teachings of the cited references, if combined, do not teach Appellants' invention (Br. 8).

Appellants argue that even if the references did teach all of the limitations of the claims, the references provide no motivation for the combination of ICSA with Anderson and DuMouchel, and the combination would render the references unsuitable for their respective intended purposes. In particular, it is argued that the network-based intrusion detection system (IDS) of ICSA, which monitors packets sent between networked computers, uses a completely different method of intrusion detection than the host-based methods of Anderson and DuMouchel, which monitor activity on specific components (Br. 8; Br. 11).

The Examiner responds that one skilled in the art would have been motivated to include a network-based sensor in Anderson "as directly suggested by ICSA, and to gain the many benefits of host-based [network-based?] intrusion detection systems listed by ICSA" (Ans. 12). The Examiner states that DuMouchel teaches a Bayes network for intrusion detection. The Examiner finds that ICSA teaches IDSs using signature and statistical analyses, and that it was known to integrate host-based and network-based IDSs (Ans. 13), and therefore, "combining such systems to include both host-based and network-based sources of data to be analyzed by the Anderson statistical and rule-based engine is [sic, was] within the capability of one of ordinary skill" (Ans. 14).

Appellants note that the Examiner submits that ICSA teaches that it was known to integrate the two technologies of host-based and network-

Appeal 2008-003088
Application 09/653,066

based IDS, but Appellants argue this provides no enabling disclosure of how to achieve such integration (Br. 11).

The Examiner responds that all three references teach collecting data and reporting it to an analysis engine, so it was within the level of skill in the art to send host-based and network-based data to a single analysis engine to gain the benefits of both systems (Ans. 15).

Appellants argue that combining ICSA with Anderson and DuMouchel would render them unsatisfactory for their intended purposes because Anderson and DuMouchel are configured to examine usage data (e.g., audit records) from a target host computer and such information cannot be derived from an examination of network traffic as taught in ICSA, and conversely, the system in ICSA could not derive network traffic activity from the host activity of Anderson and DuMouchel (Br. 12).

The Examiner disagrees, stating that ICSA deals with intrusion detection and the form of intrusion can be derived from examining network traffic (Ans. 15-16). The Examiner states that the combination of the three references only adds an additional source of data to be analyzed by Anderson's analysis engines (Ans. 16).

Appellants argue that ICSA provide no enabling teachings, guidance, or motivation to modify existing IDSs such as Anderson and DuMouchel (Br. 12).

The Examiner responds that the motivation is found in ICSA's statement that network-based and host-based system are beginning to be integrated and the technology to integrate was well known (Ans. 17-19).

Analysis

The statement of the rejection seems to both overly complicated and overly simplified. The rejection is overly complicated, because it is not necessary to modify Anderson to create a system that analyzes both host-based and network-based data, since only network-based data is claimed. It is not clear why the Examiner is relying on the statistical analysis component of Anderson when it is a host-based system. The rejection is oversimplified, because the rejection does not address that the Bayes network in DuMouchel operates on user command data in a host-based system and not on network-based data, and the rejection does not specifically address the motivation for applying a Bayes network to a different type of data in a different type of IDS.

To simplify the analysis, we mostly rely on ICSA and DuMouchel and refer to our statement of the relevant issue. Anderson is a host-based IDS and the statistical analysis configuration (Chap. 4) does not use a Bayes network, so it does not appear to add anything useful to ICSA or DuMouchel as to a Bayes network. However, Anderson discusses various types of host based statistics, which is useful in one step of the analysis.

ICSA describes that IDSs are either network-based or host-based (p. 14). Network based intrusion detection systems have sensors, and an analysis station that examine passing traffic on a network segment for signs of intrusion. "The analysis station can display the alarms or do additional analysis" (*id.* at 15). ICSA describes that both kinds of IDS generally use three kinds of available analysis engines: event or signature based, statistical analysis, and adaptive system (*id.*). Thus, ICSA teaches that a network-based IDS can have a statistical analysis component for analyzing relevant

features of network traffic for intrusion detection. We find that a network-based IDS performs the steps of "(a) receiving, at a first computing device, network traffic sent between at least a second computing device and a third computing device in said network; (b) examining the received network traffic for the presence of one or more relevant features; and (c) providing a summary or list of relevant features contained in the received network traffic [to an analysis engine]," as recited in claim 1. The difference between ICSA and the subject matter of claim 1 is that ICSA does not describe that the analysis engine is "a real-time Bayes network for analysis, the Bayes network including a model of normal traffic and a model of potentially harmful traffic as Bayes hypotheses."

DuMouchel describes using a Bayes network to analyze relevant features of a particular type of user data (i.e., user commands) on a host-based IDS. The Bayes network includes a model of normal activities (the profile of user's command sequences) which forms the null hypothesis, and a model of potentially harmful activities (that the T commands were drawn independently), which forms the alternative hypothesis. P. 6. DuMouchel does not teach utilizing a Bayes network to analyze network traffic data as opposed to user activity data. Nevertheless, obviousness is determined by the hypothetical person of ordinary skill in the art. The Supreme Court stated that "the [obviousness] analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ." *KSR*, 550 U.S. at 418. "A person of ordinary skill is also a person of ordinary creativity, not an automaton." *Id.* at 421.

DuMouchel discusses a "test for anomaly detection based on Bayesian hypothesis testing" (first page, "Introduction") and that control charts as shown in "Figures 1-3 could be constructed for virtually any intrusion detection indicator" (last page). This would reasonably have suggested to a person of ordinary skill in the IDS art that Bayesian hypothesis testing can be applied to the statistical analysis of any kind of intrusion detection data in any kind of IDS, whether host-based or network-based. DuMouchel discloses one example. It underestimates the level of ordinary skill in the art to assume that one skilled in the IDS art would limit the teachings of DuMouchel to its express teachings of analyzing user commands on a host-based IDS. For example, one of ordinary skill in the IDS art would have been motivated to apply the Bayes network teaching of DuMouchel to analyze other kinds of statistics in a host-based IDS, such as the host-based statistics discussed in Chapter 4 of Anderson, because these are just different kinds of host-based statistics. Network-based IDSs such as ICSA analyze statistics of network activity. One of ordinary skill in the IDS art would have been motivated to use any known type of statistical analysis to analyze network statistics in a network-based IDS and, therefore, would have been motivated to use the Bayes network method taught in DuMouchel, because it was a known method for analyzing statistical data in an IDS environment. One of ordinary skill is presumed to have the necessary skill to select the hypotheses for a Bayes network as applied to network traffic. DuMouchel teaches that the Bayes network had a low false alarm rate and a high rate of intrusion detection (p. 1, Summary), which advantages provide motivation for one skilled in the art to use a Bayes network. The fact that the IDS statistics are derived from network activity rather than user activity would

Appeal 2008-003088
Application 09/653,066

not have discouraged the modification, because the source of the statistics does not affect the analysis method.

Conclusion

One of ordinary skill in the IDS art would have been motivated to apply a real-time Bayes network, as taught by DuMouchel, to analysis of statistical network traffic, as taught by ICSA. The rejection of claims 1, 2, 22, and 23 is affirmed.

Claim 24

Appellants argue that Anderson, ICSA, and DuMouchel do not teach, show, or suggest all of the limitations of parent claim 23 and therefore, claim 24 is patentable for the same reasons (Br. 13).

We affirmed the rejection of parent claim 23, so this argument is not persuasive.

Appellants argue that the references do not teach the combination of parent claim 23 with dependent claim 24 (Br. 13).

The Examiner addressed claim 24 (Final Rej. 8). Appellants do not argue how the Examiner erred in either his findings of fact or conclusions of obviousness. Absent argument as to why the Examiner erred, the rejection must be affirmed. Accordingly, the rejection of claim 24 is affirmed.

Anderson, ICSA, DuMouchel, and Decasper (and others)

Claim 3

Appellants argue:

Anderson, ICSA, DuMouchel and Decasper, singly or in any permissible combination, fail to teach, show or suggest a network intrusion detection system that provides a summary or list of relevant

features of network traffic (i.e., packets or data traffic sent between network users external to a location where intrusion detection occurs) to a Bayes network for analysis, as claimed by the Appellant.

Br. 14. It is argued that Decasper does not bridge this gap (*id.* at 15).

We determined in connection with claims 1, 2, 22, and 23, that it would have been obvious to a person of ordinary skill in the IDS art to apply the Bayes network of DuMouchel to analyze a summary or list of relevant features of network traffic in a network-based IDS as taught by ICSA.

Appellants do not argue the merits of the Examiner's finding that Decasper teaches "separating the received network traffic into a plurality of sessions" or the Examiner's conclusion that it would have been obvious to modify the combination of references to provide this limitation. Arguments not made are waived. *See* 37 C.F.R. § 41.37(c)(1)(vii) ("Any arguments or authorities not included in the brief or a reply brief . . . will be refused consideration by the Board, unless good cause is shown."); *In re Watts*, 354 F.3d 1362, 1367 (Fed. Cir. 2004) ("Just as it is important that the PTO in general be barred from raising new arguments on appeal to justify or support a decision of the Board, it is important that the applicant challenging a decision not be permitted to raise arguments on appeal that were not presented to the Board." (Footnote omitted.)). *Cf. In re Baxter Travenol Labs.*, 952 F.2d 388, 391 (Fed. Cir. 1991) ("It is not the function of this court to examine the claims in greater detail than argued by an appellant, looking for nonobvious distinctions over the prior art."). There may be very good reasons why an applicant decides not to argue a limitation, e.g., the limitation may be considered well known in the art or applicant may not want patentability to be based on the particular limitation. Based on the

Appeal 2008-003088
Application 09/653,066

arguments presented, the rejection of claim 3 is affirmed for the reasons stated with respect to claims 1, 2, 22, and 23.

Claims 4-20

Appellants provide a separate heading and argument for each of dependent claims 4-20. Each argument follows the form of: (1) arguing that the combination of Anderson, ICSA, DuMouchel, and Decasper do not teach, show, or suggest all of the limitations of independent claim 3 from which the dependent claim depends and therefore, that the dependent claim is patentable for the same reasons; and (2) arguing that the combination of Anderson, ICSA, DuMouchel, and Decasper, and other references as to some claims, do not teach the limitations of claim 3 together with added limitation of the particular dependent claim.

As discussed with respect to claim 3, we conclude that it would have been obvious to a person of ordinary skill in the IDS art to apply the Bayes network of DuMouchel to analyze a summary or list of relevant features of network traffic in a network-based IDS as taught by ICSA for the reasons stated in the analysis of claims 1, 2, 22, and 23. Accordingly, the argument that the dependent claims are patentable because they depend on independent claim 3 is not persuasive.

As to the limitations of the dependent claims, Appellants argue that the references do not teach the combination of claim 3 with the particular dependent claim. However, the Examiner has addressed each of the dependent claims and Appellants do not argue how the Examiner erred in either his findings of fact or conclusions of obviousness. Absent argument as to why the Examiner erred, the rejections must be affirmed. *See*

Appeal 2008-003088
Application 09/653,066

37 C.F.R. § 41.37(c)(1)(vii) ("A statement which merely points out what a claim recites will not be considered an argument for separate patentability of the claim."); *Ernst Haas Studio, Inc. v. Palm Press, Inc.*, 164 F.3d 110, 112 (2d Cir. 1999) (declining "invitation to the court to scour the record, research any legal theory that comes to mind, and serve generally as an advocate for appellant"). Accordingly, the rejections of claims 4-20 are affirmed.

CONCLUSION

The rejections of claims 1-24 are affirmed.

Requests for extensions of time are governed by 37 C.F.R. § 1.136(b).

See 37 C.F.R. § 41.50(f).

AFFIRMED

erc

Wall & Tong, LLP
SRI International
595 Shrewsbury Avenue
Shrewsbury, NJ 07702.